Approximation errors of online sparsification criteria

Paul Honeine, *Member IEEE*

arXiv:1409.6046v1 [stat.ML] 21 Sep 2014

***Abstract*—Many machine learning frameworks, such as resource-allocating networks, kernel-based methods, Gaussian processes, and radial-basis-function networks, require a sparsification scheme in order to address the online learning paradigm. For this purpose, several online sparsification criteria have been proposed to restrict the model definition on a subset of samples. The most known criterion is the (linear) approximation criterion, which discards any sample that can be well represented by the already contributing samples, an operation with excessive computational complexity. Several computationally efficient sparsification criteria have been introduced in the literature, such as the distance, the coherence and the Babel criteria. In this paper, we provide a framework that connects these sparsification criteria to the issue of approximating samples, by deriving theoretical bounds on the approximation errors. Moreover, we investigate the error of approximating any feature, by proposing upper-bounds on the approximation error for each of the aforementioned sparsification criteria. Two classes of features are described in detail, the empirical mean and the principal axes in the kernel principal component analysis.**

***Index Terms*—Sparse approximation, adaptive filtering, kernel-based methods, resource-allocating networks, Gaussian processes, Gram matrix, machine learning, pattern recognition, online learning, sparsification criteria.**

I. Introduction

DATA DELUGE in the era of "Big Data" brings new challenges (and opportunities) in the area of machine learning and signal processing [1], [2], [3]. Demanding online learning, this paradigm cannot be addressed directly by most (if not all) conventional learning machines, such as resource-allocating networks [4], kernel-based methods for classification and regression [5], Gaussian processes [6], radial-basis-function networks [7] and kernel principal component analysis [8], only to name a few. Indeed, these machines share essentially the same underlying model, with as many parameters to be estimated as training samples, as defined by the "Representer Theorem" [9]. This model is inappropriate in online learning, where a new sample is available at each instant. To stay computationally tractable, one needs to restrict the incrementation in the model complexity, by selecting the subset of samples that contributes to a reduced-order model as an approximation of the full-order feature to be estimated.

In order to overcome this bottleneck in online learning, sparsification schemes have been proposed for all the aforementioned machines, defined as follows: at each instant, it determines if the new sample can be safely discarded from contributing to the order growth of the model; otherwise, the sample needs to take part in the order incrementation. The most known online sparsification criteria is the approximation criterion, also called approximate linear dependency. It has been widely investigated in the literature, for Gaussian processes [10], kernel recursive least squares algorithm [11], kernel least mean square algorithm [12], and kernel principal component analysis [8]. This criterion determines the relevance of discarding or accepting the current sample by comparing, to a predefined threshold, the residual error of approximating it with a representation (*i.e.*, linear combination) of samples — or nonlinearly mapped samples as in kernel methods — already contributing to the model. A crucial issue in the approximation criterion is its computational complexity, which scales cubically with the model order.

Several computationally efficient sparsification criteria have been introduced in the literature, with essentially the same computational complexity that scales linearly with the model order. These sparsification criteria rely on the topology of the samples in order to select the most relevant samples. The most widely investigated criteria are the distance and the coherence criteria, as well as the Babel criterion. The distance criterion, introduced by Platt in [4] to control the complexity of resource-allocating networks in radial-basis-function networks, retains the most mutually distant samples ; see also [13], [14] for recent advances. The coherence criterion, introduced by Honeine, Richard, and Bermudez in [15], [16] with the recent advances in compressed sensing [17], [18], retains samples that are mutually least coherent. As an extension of the coherence criterion, the Babel criterion uses the cumulative coherence as a measure of diversity [19].

These sparsification criteria have been separately investigated in the literature. To the best of our knowledge, there is no work that studies all these sparsification criteria together. The conducted analyses have been often based on the computational complexity, as advocated in [16], [20] by criticizing the computational cost of the approximation criterion in favor of the other sparsification criteria. In [15], [16], [21], we have developed with colleagues several theoretical results that allows to compare the coherence to the approximation criterion. These results have not been extended to other sparsification criteria, and were demonstrated for the particular case of unit-norm data.

This paper presents a framework to study online sparsification criteria by cross-fertilizing previously derived results, by obtaining often tighter bounds, and by extending these results to other sparsification criteria such as the distance and the Babel criteria. One the one hand, we bridge the gap between the approximation criterion and the other online sparsification criteria, firstly by providing upper bounds on the error of approximating, with samples already retained, any sample discarded by the sparsification criterion, and secondly by providing lower bounds on the error of approximating accepted samples. On the other hand, we examine the relevance of approximating any feature with a sparse model obtained with

P. Honeine is with the Institut Charles Delaunay (CNRS), Université de technologie de Troyes, 10000, Troyes, France. Phone: +33(0)325715625; Fax: +33(0)325715699; E-mail: paul.honeine@utt.fr

	Distance	Approximation	Coherence	Babel	Section
Reference: most known work	[4]	[10]	[16]	[17]	
Reference: more recent work	[20]	[8]	[25]	[19]	
Approximation of any sample	✓	·	✓	✓	**IV**
↳ Error on discarded samples	✓	·	[15]	✓	IV-A
↳ Error on any atom	✓	·	[16]	[19]	IV-B
Approximation of any feature	✓	✓	✓	✓	**V**
↳ Error on the mean (centroid)	✓	✓	[21]	✓	V-A
↳ Error on the principal axes	✓	[11]	[15]	✓	V-B

TABLE I
A BIRDS EYE VIEW OF THIS PAPER. SOME OF THE RESULTS WERE PREVIOUSLY STUDIED FOR UNIT-NORM KERNELS, AS SHOWN WITH THE REFERENCES GIVEN IN THE TABLE (WHERE · DENOTES TRIVIALITY). IN THIS WORK, WE PROVIDE AN EXTENSIVE STUDY THAT COMPLETES THE ANALYSIS TO ALL SPARSIFICATION CRITERIA, OFTEN WITH TIGHTER BOUNDS (SHOWN IN GRAY COLOR), AND WE DERIVE NEW THEORETICAL RESULTS. MOREOVER, WE GENERALIZE THESE RESULTS TO ANY TYPE OF KERNEL, BEYOND UNIT-NORM KERNELS.

any of the aforementioned sparsification criteria, including the approximation criterion. We provide upper bounds on the error of approximating any feature in the general case. Furthermore, we explore in detail two particular features, the empirical mean (*i.e.*, centroid, studied for instance in [22], [23]) and the principal axes in the kernel principal component analysis (kernel-PCA, [24]). The big picture of the cross-fertilization and extensions given in this paper is illustrated in TABLE I.

The remainder of this paper is organized as follows. Next section introduces the kernel-based machines for online learning and presents the key issues studied in this work. Section III presents the aforementioned computationally efficient sparsification criteria. Section IV investigates bounds on the error of approximation samples, either discarded or accepted by any sparsification criterion. These results are extended in Section V to the problem of approximating any feature. Section VI concludes this document with some discussions.

II. KERNEL-BASED MACHINES FOR ONLINE LEARNING

In this section, we introduce the kernel-based machines for online learning, by presenting the approximation criterion with the key issues studied in this paper.

A. Machine learning and online learning

Machine learning seeks a feature $\psi(\cdot)$ connecting an input space $\mathbb{X} \subset \mathbb{R}^d$ to an output space $\mathbb{Y} \subset \mathbb{R}$, by using a set of training samples, denoted $\{(\boldsymbol{x}_1, y_1), (\boldsymbol{x}_2, y_2), \ldots, (\boldsymbol{x}_n, y_n)\}$ with $(\boldsymbol{x}_k, y_k) \in \mathbb{X} \times \mathbb{Y}$. Considering a loss function $\mathcal{C}(\cdot, \cdot)$ defined on $\mathbb{Y} \times \mathbb{Y}$ that measures the error between the desired output and the estimated one with $\psi(\cdot)$, the optimization problem consists in minimizing a regularized empirical risk of the form

$$\underset{\psi(\cdot)\in\mathbb{H}}{\operatorname{argmin}} \sum_{i=1}^{n} \mathcal{C}(\psi(\boldsymbol{x}_i), y_i) + \eta\, \mathcal{R}(\|\psi(\cdot)\|_{\mathbb{H}}^2), \qquad (1)$$

where $\mathbb{H}$ is the feature space of candidate solutions and η is a parameter that controls the tradeoff between the fitness error (first term) and the regularity of the solution (second term) with $\mathcal{R}(\cdot)$ being a monotonically increasing function. Examples of loss functions are the quadratic loss $|\psi(\boldsymbol{x}_i) - y_i|^2$, the hinge loss $(1 - \psi(\boldsymbol{x}_i) y_i)_+$ of the SVM [5], the logistic regression $\log(1 + \exp(-\psi(\boldsymbol{x}_i) y_i))$, as well as the unsupervised loss function $-|\psi(\boldsymbol{x}_i)|^2$ which is related to the PCA.

Let $\kappa : \mathbb{X} \times \mathbb{X} \to \mathbb{R}$ be a positive definite kernel, and $(\mathbb{H}, \langle \cdot, \cdot \rangle_{\mathbb{H}})$ the induced reproducing kernel Hilbert space (RKHS) with its inner product. The reproducing property states that any function $\psi(\cdot)$ of $\mathbb{H}$ can be evaluated at any sample $\boldsymbol{x}_i$ of $\mathbb{X}$ using $\psi(\boldsymbol{x}_i) = \langle \psi(\cdot), \kappa(\cdot, \boldsymbol{x}_i) \rangle_{\mathbb{H}}$. This property shows that any sample $\boldsymbol{x}_i$ of $\mathbb{X}$ is represented with $\kappa(\cdot, \boldsymbol{x}_i)$ in the space $\mathbb{H}$. Moreover, the reproducing property leads to the so-called kernel trick, that is for any pair of samples $(\boldsymbol{x}_i, \boldsymbol{x}_j)$, we have $\langle \kappa(\cdot, \boldsymbol{x}_i), \kappa(\cdot, \boldsymbol{x}_j) \rangle_{\mathbb{H}} = \kappa(\boldsymbol{x}_i, \boldsymbol{x}_j)$. In particular, $\|\kappa(\cdot, \boldsymbol{x}_i)\|_{\mathbb{H}} = \langle \kappa(\cdot, \boldsymbol{x}_i), \kappa(\cdot, \boldsymbol{x}_i) \rangle_{\mathbb{H}} = \kappa(\boldsymbol{x}_i, \boldsymbol{x}_i)$. The most used kernels and there expressions are as follows:

Kernel	$\kappa(\boldsymbol{x}_i, \boldsymbol{x}_j)$
Linear	$\langle \boldsymbol{x}_i, \boldsymbol{x}_j \rangle$
Polynomial	$(\langle \boldsymbol{x}_i, \boldsymbol{x}_j \rangle + c)^p$
Exponential	$\exp(\langle \boldsymbol{x}_i, \boldsymbol{x}_j \rangle)$
Gaussian	$\exp\left(\frac{-1}{2\sigma^2} \|\boldsymbol{x}_i - \boldsymbol{x}_j\|^2\right)$

Among these kernels, only the Gaussian kernel is unit-norm, that is $\|\kappa(\boldsymbol{x}, \cdot)\|_{\mathbb{H}} = 1$ for any sample $\boldsymbol{x} \in \mathbb{X}$. Other kernels can be unit-norm on some restricted $\mathbb{X}$, such as the linear kernel when dealing with unit-norm samples. In this paper, we do not restrict ourselves to any particular kernel or space $\mathbb{X}$. We denote

$$r^2 = \inf_{\boldsymbol{x}\in\mathbb{X}} \kappa(\boldsymbol{x}, \boldsymbol{x}) \qquad \text{and} \qquad R^2 = \sup_{\boldsymbol{x}\in\mathbb{X}} \kappa(\boldsymbol{x}, \boldsymbol{x}) \ .$$

For unit-norm kernels, we get $R = r = 1$.

The Representer Theorem provides a principal result that is essential in kernel-based machines for classification and regression, as well as unsupervised learning. It states that the solution of the optimization problem (1) takes the form

$$\psi(\cdot) = \sum_{i=1}^{n} \alpha_i\, \kappa(\boldsymbol{x}_i, \cdot). \qquad (2)$$

The proof of this theorem is derived in [9], and a sketch of proof is given in the footnote[1]. This theorem shows that the optimal solution has as many parameters α_i to be estimated as the number of available samples $(\boldsymbol{x}_i, y_i)$. This result constitutes the principal bottleneck for online learning. Indeed, in an online setting, the solution should be adapted based on a new sample available at each instant, namely $(\boldsymbol{x}_t, y_t)$ at instant t. Thus, by including the new pair $(\boldsymbol{x}_t, y_t)$ in the training

[1]To prove the Representer Theorem (2), we decompose any function $\psi(\cdot)$ of $\mathbb{H}$ into $\psi(\cdot) = \sum_{i=1}^{n} \alpha_i\, \kappa(\boldsymbol{x}_i, \cdot) + \psi^{\perp}(\cdot)$, where $\langle \psi^{\perp}(\cdot), \kappa(\boldsymbol{x}_i, \cdot) \rangle_{\mathbb{H}} = 0$ for all $i = 1, 2, \ldots, n$. On the one hand, any evaluation $\psi(\boldsymbol{x}_i)$ is independent of $\psi^{\perp}(\cdot)$ since $\psi(\boldsymbol{x}_i) = \langle \psi(\cdot), \kappa(\boldsymbol{x}_i, \cdot) \rangle_{\mathbb{H}}$. On the other hand, the monotonically increasing function $\mathcal{R}(\cdot)$ guarantees that $\mathcal{R}(\|\psi(\cdot)\|_{\mathbb{H}}^2) = \mathcal{R}(\|\sum_{i=1}^{n} \alpha_i\, \kappa(\boldsymbol{x}_i, \cdot) + \psi^{\perp}(\cdot)\|_{\mathbb{H}}^2) \geq \mathcal{R}(\|\sum_{i=1}^{n} \alpha_i\, \kappa(\boldsymbol{x}_i, \cdot)\|_{\mathbb{H}}^2)$, where the Pythagorean theorem is used. Therefore, a null $\psi^{\perp}(\cdot)$ minimizes the regularization term without affecting the fitness term.

set, the corresponding parameter α_t is be added to the set of parameters to be estimated, by following the Representer Theorem. As a consequence, the order of the model (2) is continuously increasing.

To overcome this bottleneck, one needs to control the growth of the model order at each instant, by keeping only a fraction of the kernel functions in the expansion (2). The reduced-order model takes the form

$$\psi(\cdot) = \sum_{j=1}^{m} \alpha_j \, \kappa(\dot{\boldsymbol{x}}_j, \cdot) \tag{3}$$

with $m \ll t$, predefined or dependent on t. In this expression, $\{\dot{\boldsymbol{x}}_1, \dot{\boldsymbol{x}}_2, \ldots, \dot{\boldsymbol{x}}_m\}$ is a subset of $\{\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_t\}$, namely $\dot{\boldsymbol{x}}_j$ is some $\boldsymbol{x}_{\omega_j}$ with $\omega_j \in \{1, 2, \ldots, t\}$. We denote by dictionary the set $\mathcal{D} = \{\kappa(\dot{\boldsymbol{x}}_1, \cdot), \kappa(\dot{\boldsymbol{x}}_2, \cdot), \ldots, \kappa(\dot{\boldsymbol{x}}_m, \cdot)\}$, and by atoms its elements. Throughout this paper, all quantities associated to the dictionary have an accent (by analogy to phonetics, where stress accents are associated to prominence). This is the case for instance of the m-by-m Gram matrix $\dot{\boldsymbol{K}}$ whose (i,j)-th entry is $\kappa(\dot{\boldsymbol{x}}_i, \dot{\boldsymbol{x}}_j)$. The eigenvalues of this matrix are denoted $\dot{\lambda}_1, \dot{\lambda}_2, \ldots, \dot{\lambda}_m$, given in non-increasing order.

The optimization problem is two-fold at each instant: selecting the proper dictionary $\mathcal{D} = \{\kappa(\dot{\boldsymbol{x}}_1, \cdot), \kappa(\dot{\boldsymbol{x}}_2, \cdot), \ldots, \kappa(\dot{\boldsymbol{x}}_m, \cdot)\}$ and estimating the corresponding parameters $\alpha_1, \alpha_2, \ldots, \alpha_m$. New challenges (and opportunities) arise in an online learning setting. Determining the optimal dictionary at each instant is a combinatorial optimization problem, when optimality is measured by comparing reduced-order solution (3) to the feature in its full-order form (2). An elegant way to overcome this computationally intractable problem, is a recursive update, by determining if the new kernel function $\kappa(\boldsymbol{x}_t, \cdot)$ needs to be included to the dictionary, or it can be discarded since it is efficiently approximated with atoms already belonging to the dictionary. This is the essence of the approximation criterion.

B. Approximation criterion

The (linear) approximation criterion was initially proposed in [26] for classification and regression, and in [27] for Gaussian processes. In online learning with kernels, as studied for system identification in [11] and more recently for kernel principal component analysis in [8], it operates as follows: the current sample is discarded (not included in the dictionary), if it can be sufficiently represented by a linear combination of atoms already belonging to the dictionary; otherwise, it is included in the dictionary. Formally, the kernel function $\kappa(\boldsymbol{x}_t, \cdot)$ is included in the dictionary if

$$\min_{\xi_1 \cdots \xi_m} \left\| \kappa(\boldsymbol{x}_t, \cdot) - \sum_{j=1}^{m} \xi_j \, \kappa(\dot{\boldsymbol{x}}_j, \cdot) \right\|_{\mathrm{H}}^2 \geq \delta^2, \tag{4}$$

where δ is a positive threshold parameter that controls the level of sparseness. The above norm is the residual error obtained by projecting $\kappa(\boldsymbol{x}_t, \cdot)$ onto the space spanned by the dictionary. The optimal value of each coefficient ξ_j is obtained by nullifying the derivative of the above cost function with respect to it, which leads to

$$\boldsymbol{\xi} = \dot{\boldsymbol{K}}^{-1} \dot{\boldsymbol{\kappa}}(\dot{\boldsymbol{x}}_t),$$

where $\dot{\boldsymbol{\kappa}}(\boldsymbol{x}_t)$ is the column vector of entries $\kappa(\dot{\boldsymbol{x}}_j, \boldsymbol{x}_t)$, for $j = 1, 2, \ldots, m$. By injecting this expression in the condition (4), we get the following condition of accepting the current kernel function:

$$\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \dot{\boldsymbol{\kappa}}(\boldsymbol{x}_t)^\top \dot{\boldsymbol{K}}^{-1} \dot{\boldsymbol{\kappa}}(\boldsymbol{x}_t) \geq \delta^2. \tag{5}$$

The resulting dictionary is called δ-approximate, satisfying the relation

$$\min_{i=1 \cdots m} \min_{\xi_1 \cdots \xi_m} \left\| \kappa(\dot{\boldsymbol{x}}_i, \cdot) - \sum_{\substack{j=1 \\ j \neq i}}^{m} \xi_j \, \kappa(\dot{\boldsymbol{x}}_j, \cdot) \right\|_{\mathrm{H}} \geq \delta.$$

One could also include a removal process, in the same spirit as the fixed-budget concept, by discarding the atom that can be well approximated with the other atoms, as investigated for instance in [28]. Nonetheless, the dictionary is still δ-approximate. The use of a removal process does not affect the results given in this paper.

C. Issues studied in this paper

In the following, we describe several issues that motivates (and structures) this work, illustrated here with respect to the approximation criterion.

Computational complexity

The approximation criterion requires the inversion of the Gram matrix associated to the dictionary, which is the most computational expensive process. Its computational complexity scales cubically with the size of the dictionary, *i.e.*, $\mathcal{O}(m^3)$ operations. Moreover, the evaluation of the condition expressed in (5) requires two matrix multiplications at each instant. These computation cost may counteract the benefits of several online learning techniques, such as gradient-based and least-mean-square algorithms (*e.g.*, LMS, NLMS, affine projection, ...).

To reduce the computational burden of the approximation criterion, several computationally efficient sparsification criteria have been proposed in the literature, sharing essentially the same computational complexity that scales linearly with the size of the dictionary, *i.e.*, $\mathcal{O}(m)$ operations at each instant. The most known criteria are the distance, the coherence and the Babel criteria; see Section III for a description.

Approximation error of any sample

The approximation criterion relies on establishing a dictionary such that the error of approximating each of its atoms, with a linear combination of the other atoms, cannot be smaller than the given threshold δ. Moreover, the decision of discarding any sample from the dictionary is defined by the same process, namely when its approximation error, with a linear combination of the other atoms, is smaller than the same threshold δ. While the approximation criterion possesses such duality between accepting and discarding samples at the very same value of thresholding the approximation error, this is not the case of the other sparsification criteria.

In Section IV, we bridge the gap between the approximation criterion and the other online sparsification criteria. For this purpose, on the one hand in Section IV-A, we derive upper bounds on the error of approximating a discarded samples with atoms of a dictionary obtained by the distance, the coherence, or the Babel criterion. One the other hand in Section IV-B, we derive lower bounds on the error of approximating any atom with the other atoms of the sparse dictionary under scrutiny.

From approximating samples to approximating features

All the aforementioned sparsification criteria operate in a pre-processing scheme, by selecting samples independently of the resulting sparse representation of the feature. In other words, the selection of the relevant subset $\{\dot{\boldsymbol{x}}_1, \dot{\boldsymbol{x}}_2, \dots, \dot{\boldsymbol{x}}_m\}$ from the set $\{\boldsymbol{x}_1, \boldsymbol{x}_2, \dots, \boldsymbol{x}_t\}$ is only based on the topology of the samples; it is independent of the power of the dictionary to approximate accurately any feature of the form (2) with the reduced-order model (3).

In Section V, we study the relevance of approximating any feature with a sparse dictionary obtained by any sparsification criterion, including the approximation criterion. We derive upper bounds on the approximation error of any feature, before examining in detail two particular class of features, the empirical mean studied in Section V-A and the most relevant principal axes in kernel-PCA investigated in Section V-B.

III. Online sparsification criteria

With a novel sample $\boldsymbol{x}_t$ available at each instant t, a sparsification rule determines if $\kappa(\boldsymbol{x}_t, \cdot)$ should be included in the dictionary, by incrementing the model order m and setting $\dot{\boldsymbol{x}}_{m+1} = \boldsymbol{x}_t$. The sparsification criteria measure the relevance of such complexity-incrementation by comparing the current kernel function $\kappa(\boldsymbol{x}_t, \cdot)$ with the atoms of the dictionary. They are defined by either a dissimilarity measure, *i.e.*, constructing the dictionary with the most mutually distant atoms, or a similarity measure, *i.e.*, constructing the dictionary with the least coherent or correlated atoms. To this end, a threshold is used to control the level of sparsity of the dictionary. The most investigated criteria are outlined in the following.

A. Distance criterion

It is natural to propose a sparsification criterion that constructs a dictionary with large distances between its entries, thus discarding samples that are too close to any of the atoms already belonging to the dictionary. The current kernel function $\kappa(\boldsymbol{x}_t, \cdot)$ is included in the dictionary if

$$\min_{j=1\cdots m} \min_{\xi} \|\kappa(\boldsymbol{x}_t, \cdot) - \xi\, \kappa(\dot{\boldsymbol{x}}_j, \cdot)\|_{\mathbb{H}} \geq \delta, \tag{6}$$

for a predefined positive threshold δ; otherwise, it can be efficiently approximated, up to a multiplicative constant, with an atom of the dictionary. It is easy to see that the optimal value of the scaling factor ξ is $\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j)/\kappa(\dot{\boldsymbol{x}}_j, \dot{\boldsymbol{x}}_j)$, since the left-hand-side in the above expression is residual error of the projection of $\kappa(\boldsymbol{x}_t, \cdot)$ onto $\kappa(\dot{\boldsymbol{x}}_j, \cdot)$ (in the same spirit as the approximation criterion). This allows to simplify the condition (6) to get

$$\min_{j=1\cdots m} \left(\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \frac{\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j)^2}{\kappa(\dot{\boldsymbol{x}}_j, \dot{\boldsymbol{x}}_j)} \right) \geq \delta^2. \tag{7}$$

The resulting dictionary, called δ-distant, satisfies for any pair $(\dot{\boldsymbol{x}}_i, \dot{\boldsymbol{x}}_j)$:

$$\kappa(\dot{\boldsymbol{x}}_i, \dot{\boldsymbol{x}}_i) - \frac{\kappa(\dot{\boldsymbol{x}}_i, \dot{\boldsymbol{x}}_j)^2}{\kappa(\dot{\boldsymbol{x}}_j, \dot{\boldsymbol{x}}_j)} \geq \delta^2. \tag{8}$$

For unit-norm atoms, this expression reduces to the condition $|\kappa(\dot{\boldsymbol{x}}_i, \dot{\boldsymbol{x}}_j)| \leq \sqrt{1-\delta^2}$. This sparsification criterion has been extensively used in the literature under different names, such as the novelty criterion proposed in [4] (where the scaling factor was dropped and a prediction error mechanism was included in a second stage; see also [29], [7]) and the quantized criterion defined in [30].

B. Coherence criterion

The coherence measure has been extensively studied in the literature of compressed sensing in the particular case of the linear kernel with unit-norm samples [17], [18]. In the more general case with the kernel formalism, the coherence of a dictionary is defined with the measure

$$\max_{\substack{i,j=1\cdots m \\ i\neq j}} \frac{|\kappa(\dot{\boldsymbol{x}}_i, \dot{\boldsymbol{x}}_j)|}{\sqrt{\kappa(\dot{\boldsymbol{x}}_i, \dot{\boldsymbol{x}}_i)\, \kappa(\dot{\boldsymbol{x}}_j, \dot{\boldsymbol{x}}_j)}}, \tag{9}$$

which corresponds to the largest value of the cosine angle between any pair of atoms, since the above objective function can be written as

$$\frac{|\langle \kappa(\dot{\boldsymbol{x}}_i, \cdot), \kappa(\dot{\boldsymbol{x}}_j, \cdot) \rangle_{\mathbb{H}}|}{\|\kappa(\dot{\boldsymbol{x}}_i, \cdot)\|_{\mathbb{H}} \|\kappa(\dot{\boldsymbol{x}}_j, \cdot)\|_{\mathbb{H}}}.$$

The coherence criterion introduced in [15], [16] constructs a dictionary with atoms that are mutually least coherent, by restricting this measure below some predefined value $\gamma \in [0\,;1]$, where a null value yields an orthogonal basis. The criterion includes the current kernel function $\kappa(\boldsymbol{x}_t, \cdot)$ in the dictionary if

$$\max_{j=1\cdots m} \frac{|\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j)|}{\sqrt{\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t)\, \kappa(\dot{\boldsymbol{x}}_j, \dot{\boldsymbol{x}}_j)}} \leq \gamma. \tag{10}$$

It is worth noting that the denominator in each of the above expressions reduces to 1 when dealing with unit-norm atoms, thus expression (10) becomes

$$\max_{j=1\cdots m} |\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j)| \leq \gamma.$$

C. Babel criterion

While the coherence measure examines the largest correlation between all pairs of atoms in a dictionary, a more thorough analysis is provided by the Babel measure, which considering the maximum cumulative correlation between an atom and all the atoms of the dictionary [31], [17]. The Babel criterion for online sparsification is defined as follows: the current kernel function $\kappa(\boldsymbol{x}_t, \cdot)$ is included in the dictionary if

$$\sum_{j=1}^{m} |\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j)| \leq \gamma, \tag{11}$$

for a given positive threshold γ [19]. The resulting dictionary, called γ-Babel, satisfies

$$\max_{i=1\cdots m} \sum_{\substack{j=1 \\ j\neq i}}^{m} |\kappa(\dot{\boldsymbol{x}}_i, \dot{\boldsymbol{x}}_j)| \leq \gamma. \tag{12}$$

By analogy with the coherence measure, which corresponds to the ∞-norm when dealing with unit-norm atoms, the Babel measure[2] is related to the 1-norm of the Gram matrix, where $\|\dot{\boldsymbol{K}}\|_1 = \max_i \sum_j |\kappa(\dot{\boldsymbol{x}}_i, \dot{\boldsymbol{x}}_j)|$.

IV. Approximation error of any sample

In this section, we study the elementary issue of approximating a sample by the span of a sparse dictionary, in the kernel-based framework. To this end, this issue is considered in its two folds: one the one hand, the error of approximating a discarded sample, and on the other hand, the error of approximating any accepted sample, namely approximating any atom of the dictionary with all the other atoms. We provide upper bounds on the former and lower bounds on the latter, for each of the sparsification criteria studied in previous section. It is worth noting that only the approximation criterion relies on a duality of discarding and accepting samples at the very same value in thresholding the approximation error, which is not the case of the other criteria, as examined in the following.

Let $\dot{\mathcal{P}}$ be the projection operator onto the subspace spanned by the atoms $\kappa(\dot{\boldsymbol{x}}_1, \cdot), \ldots, \kappa(\dot{\boldsymbol{x}}_m, \cdot)$ of a dictionary resulting from a sparsification criterion. Thus, for any sample $\boldsymbol{x}$, the projection of the kernel function $\kappa(\boldsymbol{x}, \cdot)$ onto this subspace is given by $\dot{\mathcal{P}}\kappa(\boldsymbol{x}, \cdot)$. The quadratic norm of the latter corresponds to the maximum inner product $\langle \kappa(\boldsymbol{x}, \cdot), \varphi(\cdot) \rangle_{\mathbb{H}}$ over all the unit-norm functions $\varphi(\cdot)$ of that subspace. By writing $\varphi(\cdot) = \sum_{j=1}^m \beta_j \kappa(\dot{\boldsymbol{x}}_j, \cdot) / \|\sum_{j=1}^m \beta_j \kappa(\dot{\boldsymbol{x}}_j, \cdot)\|_{\mathbb{H}}$, one gets

$$\begin{aligned} \|\dot{\mathcal{P}}\kappa(\boldsymbol{x}, \cdot)\|_{\mathbb{H}}^2 &= \max_{\boldsymbol{\beta}} \frac{\langle \sum_{j=1}^m \beta_j \kappa(\dot{\boldsymbol{x}}_j, \cdot), \kappa(\boldsymbol{x}, \cdot) \rangle_{\mathbb{H}}}{\|\sum_{j=1}^m \beta_j \kappa(\dot{\boldsymbol{x}}_j, \cdot)\|_{\mathbb{H}}} \\ &= \max_{\boldsymbol{\beta}} \frac{\sum_{j=1}^m \beta_j \kappa(\boldsymbol{x}, \dot{\boldsymbol{x}}_j)}{\|\sum_{j=1}^m \beta_j \kappa(\dot{\boldsymbol{x}}_j, \cdot)\|_{\mathbb{H}}}. \end{aligned} \tag{13}$$

Moreover, the Pythagorean Theorem allows to measure the residual norm of this projection, with

$$\begin{aligned} \|(\mathbf{I} - \dot{\mathcal{P}})\kappa(\boldsymbol{x}, \cdot)\|_{\mathbb{H}}^2 &= \|\kappa(\boldsymbol{x}, \cdot)\|_{\mathbb{H}}^2 - \|\dot{\mathcal{P}}\kappa(\boldsymbol{x}, \cdot)\|_{\mathbb{H}}^2 \\ &= \kappa(\boldsymbol{x}, \boldsymbol{x}) - \|\dot{\mathcal{P}}\kappa(\boldsymbol{x}, \cdot)\|_{\mathbb{H}}^2, \end{aligned}$$

where $\mathbf{I}$ is the identity operator. Therefore, the quadratic approximation error is

$$\|(\mathbf{I} - \dot{\mathcal{P}})\kappa(\boldsymbol{x}, \cdot)\|_{\mathbb{H}}^2 = \kappa(\boldsymbol{x}, \boldsymbol{x}) - \max_{\boldsymbol{\beta}} \frac{\sum_{j=1}^m \beta_j \kappa(\boldsymbol{x}, \dot{\boldsymbol{x}}_j)}{\|\sum_{j=1}^m \beta_j \kappa(\dot{\boldsymbol{x}}_j, \cdot)\|_{\mathbb{H}}}. \tag{14}$$

In the following, we investigate this expression in order to derive proper bounds on the approximation error, either when the sample is discarded or when it already belongs to the dictionary.

[2]One can also consider a normalized version of the Babel measure, by substituting $\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j)$ in (11) with $\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j) / \sqrt{\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t)\, \kappa(\dot{\boldsymbol{x}}_j, \dot{\boldsymbol{x}}_j)}$. These two definitions are equivalent when dealing with unit-norm atoms. To the best of our knowledge, this formulation is not used in the literature. Moreover, it looses the matrix-norm notion.

A. Approximation error of discarded samples

When the sample $\boldsymbol{x}_t$ is discarded, we propose to upper bound the quadratic approximation error (14) with

$$\|(\mathbf{I} - \dot{\mathcal{P}})\kappa(\boldsymbol{x}_t, \cdot)\|_{\mathbb{H}}^2 \leq \kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \max_j \frac{|\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j)|}{\sqrt{\kappa(\dot{\boldsymbol{x}}_j, \dot{\boldsymbol{x}}_j)}}, \tag{15}$$

where the inequality corresponds to the special choice of the coefficients, with $\beta_1 = \ldots = \beta_m = 0$ except for $\beta_j = \operatorname{sign}(\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j))$. Next, we show that the quotient $|\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j)| / \sqrt{\kappa(\dot{\boldsymbol{x}}_j, \dot{\boldsymbol{x}}_j)}$ in the above expression is bounded, with a lower bound that depends on the threshold of the investigated sparsification criterion. For this purpose, we examine separately the distance (Theorem 1), the coherence (Theorem 2), and the Babel (Theorem 3) criteria.

Theorem 1 (Discarding error for the distance criterion): Let $\boldsymbol{x}_t$ be a sample not satisfying the distance condition (7) for some given threshold δ. The quadratic error of approximating $\kappa(\boldsymbol{x}_t, \cdot)$ with a linear combination of atoms from the resulting dictionary is upper-bounded by

$$\delta^2 \qquad \text{and} \qquad \kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \sqrt{\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \delta^2}.$$

The latter upper bound is sharper when $\delta^2 > \kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - 1$. This is the case when dealing with unit-norm atoms, where we get the upper bound $1 - \sqrt{1 - \delta^2}$.

Proof: Firstly, one can easily derive the first expression of the upper bound, since

$$\begin{aligned} \min_{\xi_1 \cdots \xi_m} \Big\| \kappa(\boldsymbol{x}_t, \cdot) - \sum_{i=1}^m \xi_i\, \kappa(\dot{\boldsymbol{x}}_i, \cdot) \Big\|_{\mathbb{H}}^2 & \\ \leq \min_{j=1\cdots m} \min_{\xi_j} \big\| \kappa(\boldsymbol{x}_t, \cdot) - \xi_j\, \kappa(\dot{\boldsymbol{x}}_j, \cdot) \big\|_{\mathbb{H}}^2 & \\ < \delta^2, & \end{aligned}$$

where the first inequality follows from the special case when all ξ_i are null except for a single one, and the second inequality is due to the violation of (6).

Secondly, the second expression of the upper bound is a bit more trickier. The approximation error is given by the norm of the residual of the projection of $\kappa(\boldsymbol{x}_t, \cdot)$ onto the subspace spanned by the dictionary atoms, namely as given in (15). Since $\boldsymbol{x}_t$ does not satisfy the condition (6)-(7), we have

$$\min_{j=1\ldots m} \left(\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \frac{\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j)^2}{\kappa(\dot{\boldsymbol{x}}_j, \dot{\boldsymbol{x}}_j)} \right) < \delta^2,$$

and as a consequence, since $\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) \geq \delta^2$, we can easily show that

$$\sqrt{\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \delta^2} < \max_{j=1\ldots m} \frac{|\kappa(\boldsymbol{x}_t, \dot{\boldsymbol{x}}_j)|}{\sqrt{\kappa(\dot{\boldsymbol{x}}_j, \dot{\boldsymbol{x}}_j)}}.$$

By injecting this inequality in (15), we get the second expression of the upper bound with

$$\big\| \big(\mathbf{I} - \dot{\mathcal{P}}\big)\kappa(\boldsymbol{x}_t, \cdot) \big\|_{\mathbb{H}}^2 < \kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \sqrt{\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \delta^2}.$$

Finally, we compare these two expressions. It is easy to see that $\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \sqrt{\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \delta^2}$ is sharper than δ^2 when

$$\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \sqrt{\kappa(\boldsymbol{x}_t, \boldsymbol{x}_t) - \delta^2} < \delta^2,$$

namely when $\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)-\delta^2<\sqrt{\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)-\delta^2}$. This condition, of the form $u^2<u$, is satisfied when $u<1$, namely $\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)-\delta^2<1$. ∎

Theorem 2 (Discarding error for the coherence criterion): Let $\boldsymbol{x}_t$ be a sample not satisfying the coherence condition (10) for some given threshold γ. The quadratic error of approximating $\kappa(\boldsymbol{x}_t,\cdot)$ with a linear combination of atoms from the resulting dictionary is upper-bounded by

$$\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)-\gamma\sqrt{\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)}.$$

In the particular case of unit-norm atoms, we get $1-\gamma$.

Proof: The unfulfilled coherence condition (10), namely

$$\max_{j=1\cdots m}\frac{|\kappa(\boldsymbol{x}_t,\grave{\boldsymbol{x}}_j)|}{\sqrt{\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)\,\kappa(\grave{\boldsymbol{x}}_j,\grave{\boldsymbol{x}}_j)}}>\gamma,$$

can be written in the equivalent form

$$\max_{j=1\cdots m}\frac{|\kappa(\boldsymbol{x}_t,\grave{\boldsymbol{x}}_j)|}{\sqrt{\kappa(\grave{\boldsymbol{x}}_j,\grave{\boldsymbol{x}}_j)}}>\gamma\sqrt{\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)}.$$

By injecting this inequality in (15), we get an upper bound on the approximation error as follows:

$$\|(\mathbf{I}-\grave{\mathcal{P}})\kappa(\boldsymbol{x}_t,\cdot)\|^2_{\mathbb{H}}<\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)-\gamma\,\sqrt{\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)},$$

which concludes the proof. ∎

Theorem 3 (Discarding error for the Babel criterion): Let $\boldsymbol{x}_t$ be a sample not satisfying the Babel condition (11) for some given threshold γ. The quadratic error of approximating $\kappa(\boldsymbol{x}_t,\cdot)$ with a linear combination of atoms from the resulting dictionary is upper-bounded by

$$\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)-\frac{\gamma}{\sqrt{m(R^2+\gamma)}},$$

which becomes $1-\frac{\gamma}{\sqrt{m(1+\gamma)}}$ for unit-norm atoms.

Proof: To prove this result, we use the quadratic approximation error given in expression (14) where, for the particular case of $\beta_j=\operatorname{sign}(\kappa(\boldsymbol{x}_t,\grave{\boldsymbol{x}}_j))$, we get

$$\begin{aligned}\|(\mathbf{I}-\grave{\mathcal{P}})\kappa(\boldsymbol{x}_t,\cdot)\|^2_{\mathbb{H}}&=\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)-\max_{\boldsymbol{\beta}}\frac{\sum_{j=1}^m\beta_j\kappa(\boldsymbol{x}_t,\grave{\boldsymbol{x}}_j)}{\|\sum_{j=1}^m\beta_j\kappa(\grave{\boldsymbol{x}}_j,\cdot)\|_{\mathbb{H}}}\\&\leq\kappa(\boldsymbol{x}_t,\boldsymbol{x}_t)-\frac{\sum_{j=1}^m|\kappa(\boldsymbol{x}_t,\grave{\boldsymbol{x}}_j)|}{(\boldsymbol{\beta}^\top\grave{\boldsymbol{K}}\boldsymbol{\beta})^{\frac{1}{2}}}\end{aligned}$$

The above numerator is bounded since the Babel condition (11) is not satisfied, namely $\sum_{j=1}^m|\kappa(\boldsymbol{x}_t,\grave{\boldsymbol{x}}_j)|>\gamma$. The above denominator is bounded thanks to the min-max theorem (*i.e.*, the Rayleigh-Ritz quotient), with

$$\boldsymbol{\beta}^\top\grave{\boldsymbol{K}}\boldsymbol{\beta}\leq\grave{\lambda}_1\|\boldsymbol{\beta}\|^2.$$

It turns out that the this upper bound is equal to $m(R^2+\gamma)$. To show this, on the one hand, we have $\sum_{j=1}^m\beta_j^2=\sum_{j=1}^m|\operatorname{sign}(\kappa(\boldsymbol{x}_t,\grave{\boldsymbol{x}}_j))|^2=m$ due to the aforementioned particular choice of the β_j. On the other hand, the eigenvalues of a Gram matrix $\grave{\boldsymbol{K}}$ associated to a γ-Babel dictionary are upper-bounded by $R^2+\gamma$, as given in the Appendix; see [32, Theorem 5] for more details. The combination of all these results concludes the proof. ∎

B. Approximation error of an atom from the dictionary

In this section, we study the approximation of an atom of a dictionary with a linear combination of its other atoms. We provide a lower bound on the approximation error for each sparsification criterion.

Let $\kappa(\grave{\boldsymbol{x}}_i,\cdot)$ be an atom of the dictionary, and consider its projection onto the subspace spanned by the other $m-1$ atoms. By following the same derivations as in the beginning of Section IV, we have

$$\|(\mathbf{I}-\grave{\mathcal{P}})\kappa(\grave{\boldsymbol{x}}_i,\cdot)\|^2_{\mathbb{H}}=\kappa(\grave{\boldsymbol{x}}_i,\grave{\boldsymbol{x}}_i)-\max_{\boldsymbol{\beta}}\frac{\sum_{j=1,j\neq i}^m\beta_j\kappa(\grave{\boldsymbol{x}}_i,\grave{\boldsymbol{x}}_j)}{\|\sum_{j=1,j\neq i}^m\beta_j\kappa(\grave{\boldsymbol{x}}_j,\cdot)\|_{\mathbb{H}}}.$$

On the one hand, the numerator in the above expression is upper-bounded, since we have from the Cauchy-Schwarz inequality:

$$\left(\sum_{\substack{j=1\\j\neq i}}^m\beta_j\kappa(\grave{\boldsymbol{x}}_i,\grave{\boldsymbol{x}}_j)\right)^2\leq\sum_{\substack{j=1\\j\neq i}}^m\beta_j^2\;\sum_{\substack{j=1\\j\neq i}}^m|\kappa(\grave{\boldsymbol{x}}_i,\grave{\boldsymbol{x}}_j)|^2.$$

On the other hand, the denominator has a lower bound, since we have:

$$\Big\|\sum_{\substack{j=1\\j\neq i}}^m\beta_j\kappa(\grave{\boldsymbol{x}}_j,\cdot)\Big\|^2_{\mathbb{H}}=\boldsymbol{\beta}^\top\grave{\boldsymbol{K}}_{\backslash\{i\}}\boldsymbol{\beta}\geq\grave{\lambda}_{\backslash\{i\}m-1}\|\boldsymbol{\beta}\|^2,$$

where $\grave{\boldsymbol{K}}_{\backslash\{i\}}$ is the $(m-1)$-by-$(m-1)$ submatrix of the Gram matrix $\grave{\boldsymbol{K}}$ obtained by removing its i-th row and its i-th column, *i.e.*, the entries associated to $\grave{\boldsymbol{x}}_i$, and $\grave{\lambda}_{\backslash\{i\}m-1}$ is its smallest eigenvalue. By combining these two inequalities, we get

$$\|(\mathbf{I}-\grave{\mathcal{P}})\kappa(\grave{\boldsymbol{x}}_i,\cdot)\|^2_{\mathbb{H}}\geq\kappa(\grave{\boldsymbol{x}}_i,\grave{\boldsymbol{x}}_i)-\sqrt{\frac{1}{\grave{\lambda}_{\backslash\{i\}m-1}}\sum_{\substack{j=1\\j\neq i}}^m|\kappa(\grave{\boldsymbol{x}}_i,\grave{\boldsymbol{x}}_j)|^2},\tag{16}$$

This expression is investigated in the following for the distance (Theorem 4), the coherence (Theorem 5), and the Babel (Theorem 6) criteria. For each sparsification criterion, the above lower bound is written by using the corresponding summation expression and the appropriate lower bound on the eigenvalues, as derived in [32, Section IV] and put in a nutshell in the Appendix.

Theorem 4 (Acceptance error for the distance criterion): For a dictionary resulting from the distance criterion for some given threshold δ, the quadratic error of approximating any atom $\kappa(\grave{\boldsymbol{x}}_i,\cdot)$ with a linear combination of the other atoms is lower-bounded by

$$\kappa(\grave{\boldsymbol{x}}_i,\grave{\boldsymbol{x}}_i)-\sqrt{\frac{\big(\kappa(\grave{\boldsymbol{x}}_i,\grave{\boldsymbol{x}}_i)-\delta^2\big)(m-1)R^2}{r^2-(m-2)R\sqrt{R^2-\delta^2}}}.$$

For unit-norm atoms, we get a lower bound for all atoms, with

$$1-\sqrt{\frac{(m-1)(1-\delta^2)}{1-(m-2)\sqrt{1-\delta^2}}}.$$

Proof: The proof is split in two parts, by investigating expression (16). Firstly, the summation term is upper-bounded since, from (8), we have that any pair $(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)$ satisfies

$$\begin{aligned}\sum_{\substack{j=1\\ j\neq i}}^{m} |\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)|^2 &\leq \sum_{\substack{j=1\\ j\neq i}}^{m} \kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j) \left(\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) - \delta^2\right) \\ &= \left(\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) - \delta^2\right) \sum_{\substack{j=1\\ j\neq i}}^{m} \kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j) \\ &= \left(\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) - \delta^2\right)(m-1) \max_{\substack{j=1\cdots m\\ j\neq i}} \kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j) \\ &= \left(\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) - \delta^2\right)(m-1)R^2.\end{aligned}$$

Secondly, the eigenvalue in this expression is lower-bounded by $r^2-(m-2)R\sqrt{R^2-\delta^2}$ for a δ-distant dictionary of $m-1$ atoms, as shown in Lemma A.1 of the Appendix. ∎

Theorem 5 (Acceptance error for the coherence criterion): For a dictionary resulting from the coherence criterion for some given threshold γ, the quadratic error of approximating any atom $\kappa(\grave{\boldsymbol{x}}_i, \cdot)$ with a linear combination of the other atoms is lower-bounded by

$$\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) - \sqrt{\frac{(m-1)\,\gamma^2 R^2 \kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i)}{r^2-(m-2)\gamma R^2}}.$$

For unit-norm atoms, this bounds becomes independent of $\grave{\boldsymbol{x}}_i$, with

$$1 - \sqrt{\frac{(m-1)\,\gamma^2}{1-(m-2)\,\gamma}}.$$

Proof: The proof follows the same procedure as in the previous proof. On the one hand, we have

$$\begin{aligned}\sum_{\substack{j=1\\ j\neq i}}^{m} |\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)|^2 &\leq (m-1) \max_{\substack{j=1\cdots m\\ j\neq i}} |\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)|^2 \\ &\leq (m-1)\,\gamma^2 \max_{\substack{j=1\cdots m\\ j\neq i}} \kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i)\,\kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j) \\ &\leq (m-1)\,\gamma^2 R^2 \kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i),\end{aligned}$$

where the second inequality follows from the coherence condition. On the other hand, we use the lower bound $r^2-(m-2)\gamma R^2$ on the eigenvalues associated to a γ-coherent dictionary of $m-1$ atoms, as derived in Lemma A.2 of the Appendix. To complete the proof, we combine these results in (16). ∎

Theorem 6 (Acceptance error for the Babel criterion): For a dictionary resulting from the Babel criterion for some given threshold γ, the quadratic error of approximating any atom $\kappa(\grave{\boldsymbol{x}}_i, \cdot)$ with a linear combination of the other atoms is lower-bounded by

$$\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) - \frac{\gamma}{\sqrt{r^2-\gamma}}.$$

For unit-norm atoms, we get the following lower bound for all atoms:

$$1 - \frac{\gamma}{\sqrt{1-\gamma}}.$$

Proof: The proof is obtained by substituting the ℓ_2-norm in (16), *i.e.*, $\left(\sum_j |\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)|^2\right)^{\frac{1}{2}}$, with an ℓ_1-norm, since we have the relation $\|\boldsymbol{u}\|_2 \leq \|\boldsymbol{u}\|_1$. This yields

$$\|(\mathbf{I} - \grave{\mathcal{P}})\kappa(\grave{\boldsymbol{x}}_i, \cdot)\|_{\mathrm{H}}^2 \geq \kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) - \frac{1}{\sqrt{\grave{\lambda}_{\setminus\{i\}m-1}}} \sum_{\substack{j=1\\ j\neq i}}^{m} |\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)|.$$

The above summation term is upper-bounded by γ thanks to the Babel definition in (12). Moreover, the above eigenvalue is lower-bounded by $r^2-\gamma$, as derived in Lemma A.3 of the Appendix. This concludes the proof. ∎

This theorem should be compared with the work of [19, Theorem 1], where the authors propose a lower bound on the quadratic approximation error for unit-norm atoms, with $1-\gamma$. It is easy to see that the lower bound given in Theorem 6 is tighter than the previously proposed bound, and extends the result to atoms that are not unit-norm.

V. Approximation of a feature

In this section, we study the relevance of approximating any feature with its projection onto the subspace spanned by the atoms of a dictionary. An upper bound on the approximation error is derived in the following theorem for any sparsification criterion, while specific bounds in term of the threshold of each criterion are given in the following Theorem 8. Moreover, these results are explored in two particular kernel-based learning algorithms, with the empirical mean (see Section V-A) and the principal axes (see Section V-B) as features to be estimated.

Theorem 7: Consider the approximation of some feature $\psi(\cdot) = \sum_{i=1}^{n} \alpha_i\, \kappa(\boldsymbol{x}_i, \cdot)$ with a sparse solution given by projecting it onto the subspace spanned by the m atoms of a given dictionary. The quadratic error of such approximation is upper-bounded by

$$(n-m)\ \|\boldsymbol{\alpha}\|^2\, \epsilon^2,$$

where ϵ is an upper bound on the approximation of any $\kappa(\boldsymbol{x}_i, \cdot)$ with a linear combination of atoms from the dictionary.

Proof: Let $\grave{\mathcal{P}}$ be the projection operator onto the subspace spanned by the atoms of the dictionary under scrutiny, approximating $\psi(\cdot) = \sum_{i=1}^{n} \alpha_i \kappa(\boldsymbol{x}_i, \cdot)$ with $\grave{\psi}(\cdot) = \sum_{j=1}^{m} \grave{\alpha}_j \kappa(\grave{\boldsymbol{x}}_j, \cdot)$. The error of such approximation is

$$\begin{aligned}\|(\mathbf{I} - \grave{\mathcal{P}})\psi(\cdot)\|_{\mathrm{H}} &= \|\sum_{i=1}^{n} \alpha_i\, (\mathbf{I} - \grave{\mathcal{P}})\, \kappa(\boldsymbol{x}_i, \cdot)\|_{\mathrm{H}} \\ &\leq \sum_{i=1}^{n} |\alpha_i|\, \|(\mathbf{I} - \grave{\mathcal{P}})\, \kappa(\boldsymbol{x}_i, \cdot)\|_{\mathrm{H}}, \qquad (17)\end{aligned}$$

where the inequality is due to the generalized triangular inequality. By applying the Cauchy-Schwarz inequality, we get the quadratic approximation error

$$\|(\mathbf{I} - \grave{\mathcal{P}})\psi(\cdot)\|_{\mathrm{H}}^2 \leq \sum_{i=1}^{n} \alpha_i^2 \ \sum_{i=1}^{n} \|(\mathbf{I} - \grave{\mathcal{P}})\, \kappa(\boldsymbol{x}_i, \cdot)\|_{\mathrm{H}}^2. \qquad (18)$$

The first summation is the quadratic ℓ_2-norm of the vector of coefficients, namely $\|\boldsymbol{\alpha}\|^2$. For the second summation, we separate it in two terms, entries belonging to the dictionary

and those discarded thanks to the used sparsification criterion. While the former do not contribute to the error, only the latter take part in the summation, namely the $n-m$ discarded samples where m is the size of the dictionary. Let ϵ^2 be an upper bound on the quadratic error of discarding samples, as given in Section IV-A. Then, we get

$$\|(\mathbf{I}-\hat{\mathcal{P}})\psi(\cdot)\|^2_{\mathbb{H}} \leq (n-m)\,\|\boldsymbol{\alpha}\|^2\,\epsilon^2,$$

which concludes the proof. ∎

By revisiting the upper bounds given in Section IV-A for each sparsification criterion, we can easily show the following results. Expressions for non-unit-norm atoms can be derived without difficulty from Theorems 1, 2, 3.

Theorem 8: The upper bound given in Theorem 7 can be specified for each sparsification criterion in terms of the used threshold. For unit-norm atoms, we have

- $(n-m)\|\boldsymbol{\alpha}\|^2(1-\sqrt{1-\delta^2})$ for the δ-distant criterion.
- $(n-m)\|\boldsymbol{\alpha}\|^2\delta^2$ for the δ-approximate criterion.
- $(n-m)\|\boldsymbol{\alpha}\|^2(1-\gamma)$ for the γ-coherent criterion.
- $(n-m)\|\boldsymbol{\alpha}\|^2\big(1-\gamma/\sqrt{m(1+\gamma)}\big)$ for the γ-Babel criterion.

We explore next these results for two particular kernel-based learning algorithms, in order to clarify the relevance of these bounds.

A. Approximation of the empirical mean

The empirical mean is a fundamental feature of the set of sample, and its use is essential in many statistical methods. For instance, it is investigated in [23] for visualization and clustering of nonnegative data and in [22], [33] for one-class classification with kernel-based methods. In the following, we study the relevance of approximating the empirical mean by its projection onto the subspace spanned by the atoms of a dictionary. Let $\psi(\cdot)=\frac{1}{n}\sum_{i=1}^n\kappa(\boldsymbol{x}_i,\cdot)$ be the empirical mean, namely $\alpha_i=1/n$ for any $i=1,2,\ldots,n$. From Theorem 7, we get

$$\|(\mathbf{I}-\hat{\mathcal{P}})\psi(\cdot)\|^2_{\mathbb{H}} \leq \left(1-\frac{m}{n}\right)\epsilon^2, \tag{19}$$

where $\max_i\|(\mathbf{I}-\hat{\mathcal{P}})\kappa(\boldsymbol{x}_i,\cdot)\|_{\mathbb{H}}\leq\epsilon$. In the following, we give a sharper bound.

Indeed, we provide a sharper bound by relaxing the use of the Cauchy-Schwarz inequality in (18), thanks to the fact that the coefficients α_i are constant, *i.e.*, independent of i. As a consequence, we get by revisiting expression (17):

$$\begin{aligned}\|(\mathbf{I}-\hat{\mathcal{P}})\psi(\cdot)\|_{\mathbb{H}} &\leq \sum_{i=1}^n|\alpha_i|\|(\mathbf{I}-\hat{\mathcal{P}})\,\kappa(\boldsymbol{x}_i,\cdot)\|_{\mathbb{H}}\\ &= \frac{1}{n}\sum_{i=1}^n\|(\mathbf{I}-\hat{\mathcal{P}})\,\kappa(\boldsymbol{x}_i,\cdot)\|_{\mathbb{H}}\\ &\leq \frac{1}{n}\,(n-m)\,\epsilon,\end{aligned}$$

where we have followed the same decomposition as in the proof of Theorem 7, with only the $n-m$ discarded samples contribute to the summation term. Therefore, the quadratic approximation error is upper-bounded as follows:

$$\|(\mathbf{I}-\hat{\mathcal{P}})\psi(\cdot)\|^2_{\mathbb{H}} \leq \left(1-\frac{m}{n}\right)^2\epsilon^2.$$

This bound is sharper than the one in (19) since $1-\frac{m}{n}<1$.

By revisiting Theorem 8 in the light of this result, the upper bound on the quadratic approximation error $\|(\mathbf{I}-\hat{\mathcal{P}})\psi(\cdot)\|^2_{\mathbb{H}}$ can be described in terms of the threshold of each sparsification criterion, as follows:

- $\left(1-\frac{m}{n}\right)^2\left(1-\sqrt{1-\delta^2}\right)$ for the δ-distant criterion.
- $\left(1-\frac{m}{n}\right)^2\delta^2$ for the δ-approximate criterion.
- $\left(1-\frac{m}{n}\right)^2(1-\gamma)$ for the γ-coherent criterion.
- $\left(1-\frac{m}{n}\right)^2\left(1-\frac{\gamma}{\sqrt{m(1+\gamma)}}\right)$ for the γ-Babel criterion.

These results generalize the work in [21], where only the case of the coherence criterion is studied. Expressions for dictionaries with atoms that are not unit-norm can be easily obtained thanks to Theorems 1, 2, 3.

B. Approximation of the most relevant principal axes

Any sparsification criterion can be seen as a dimensionality reduction technique, because it identifies a subspace by selecting relevant samples from the available ones. Since it is an unsupervised approach, it is natural to connect it with the kernel principal component analysis. For the sake of clarity, it is assumed that the data are centered in the feature space; see [34] for connections to the uncentered case.

The principal component analysis (PCA) seeks the principal axes that capture the most of the data variance. The principal axes correspond to the eigenvectors associated to the largest eigenvalues of the covariance matrix. In its kernel-based counterpart, *i.e.*, the kernel-PCA, the k-th principal axis takes the form $\psi_k(\cdot)=\sum_{i=1}^n\alpha_{i,k}\,\kappa(\boldsymbol{x}_i,\cdot)$, where the coefficients $\alpha_{i,k}$ are the entries of the k-th eigenvector of the Gram matrix $\boldsymbol{K}$. Moreover, to get unit-norm principal axes, the coefficients $\alpha_{i,k}$ are normalized such that $\sum_{i=1}^n\alpha_{i,k}^2=1/n\lambda_k$. In this expression, λ_k is the k-th eigenvalue of the Gram matrix, also called principal value. In the following, we highlight the connections between the kernel-PCA and the online sparsification criteria.

Theorem 9: Let $\psi_k(\cdot)$ be the k-th principal axe of the kernel functions $\kappa(\boldsymbol{x}_i,\cdot)$, for $i=1,2,\ldots,n$, associated to the eigenvalue λ_k of the corresponding Gram matrix. Its approximation with a dictionary of m kernel functions has a quadratic error that can be upper-bounded by

$$\left(1-\frac{m}{n}\right)\frac{\epsilon^2}{\lambda_k},$$

where ϵ is an upper bound on the approximation of any $\kappa(\boldsymbol{x}_i,\cdot)$ with a linear combination of atoms from the dictionary.

The proof of this theorem is straightforward, by substituting $\|\boldsymbol{\alpha}\|^2$ with $1/n\lambda_k$ in Theorem 7. Theorem 9 shows that, under the only condition that the used dictionary has an upper bound on the error of approximating each kernel function, the principal axes associated to the largest principal values have the smallest approximation errors. One can therefore say that the most relevant principal axes lie, with a small error, in the span of the sparse dictionary.

Moreover, we derive expressions for each sparsification criterion, as given next in terms of the used threshold:

- $\left(1-\frac{m}{n}\right)\frac{1-\sqrt{1-\delta^2}}{\lambda_k}$ for the δ-distant criterion.
- $\left(1-\frac{m}{n}\right)\frac{\delta^2}{\lambda_k}$ for the δ-approximate criterion.
- $\left(1-\frac{m}{n}\right)\frac{1-\gamma}{\lambda_k}$ for the γ-coherent criterion.
- $\left(1-\frac{m}{n}\right)\frac{1}{\lambda_k}\left(1-\frac{\gamma}{\sqrt{m(1+\gamma)}}\right)$ for the γ-Babel criterion.

These results generalize previous work on the approximation and the coherence criteria, and provide tighter bounds than the ones previously known in the literature. Indeed, the upper bound δ^2/λ_k was derived for the approximation criterion in [11, Theorem 3.3] and in [8, Theorem 5], while the coherence criterion is studied in [15, Proposition 5] with the upper bound $(1-\gamma)/\lambda_k$.

VI. Final remarks

In this paper, we studied the approximation errors of any sample when dealing with the distance, the coherence, or the Babel criterion, revealing that these criteria are roughly based on an approximation process. By deriving an upper bound on the error of approximating a sample discarded from the dictionary, we explored that the atoms are "sufficient" to represent any sample. The dual condition, namely showing that each atom of the dictionary is "necessary", was also exhibited by providing a lower bound on the approximation of any atom of the dictionary with the other atoms. Moreover, beyond the analysis of a single sample, we extended these results to the estimation of any feature, by describing in detail two classes of features, the empirical mean (*i.e.*, centroid) and the principal axes in kernel-PCA.

This work did not devise any particular sparsification criterion. It provided a framework to study online sparsification criteria. We argued that these criteria behave essentially in an identical mechanism, and share many interesting and desirable properties. Without loss of generality, we considered the framework of kernel-based learning algorithms. It is worth noting that these machines are intimately connected with the Gaussian processes [6], where the approximation criterion was initially proposed [10].

Appendix

This appendix provides bounds on the eigenvalues of a Gram matrix associated to a sparse dictionary, for each of the sparsity measures investigated in this paper. For completeness, these bounds are put here in a nutshell ; see [32, Section IV] for more details. A cornerstone of these results is the well-known Geršgorin Discs Theorem [35, Chapter 6]. Revisited here for a Gram matrix associated to a sparse dictionary, it states that any of its eigenvalues lies in the union of the m discs, centered on each diagonal entry of $\grave{\boldsymbol{K}}$ with a radius given by the sum of the absolute values of the other $m-1$ entries from the same row. In other words, for each $\grave{\lambda}_i$, there exists at least one $j \in \{1, 2, \ldots, m\}$ such that

$$|\grave{\lambda}_i - \kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j)| \leq \sum_{\substack{j=1 \\ j\neq i}}^{m} |\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)|. \tag{20}$$

This theorem provides upper and lower bounds on the eigenvalues of a Gram matrix associated to a sparse dictionary, as described in the following for each sparsity measure.

Lemma A.1: The eigenvalues of a Gram matrix associated to a δ-distant dictionary of m atoms are bounded as follows:

$$r^2 - (m-1)R\sqrt{R^2-\delta^2} \leq \grave{\lambda}_m \leq \cdots \leq \grave{\lambda}_1 \leq R^2 + (m-1)R\sqrt{R^2-\delta^2}.$$

Proof: From (8), a δ-distant dictionary satisfies

$$|\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)| \leq \sqrt{\kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j)\left(\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) - \delta^2\right)},$$

for any $i = 1, 2, \ldots, m$, which yields

$$\begin{aligned}\sum_j |\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)| &\leq \sum_j \sqrt{\kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j)\left(\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) - \delta^2\right)} \\ &= \sqrt{\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) - \delta^2} \sum_j \sqrt{\kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j)}.\end{aligned}$$

By substituting this relation in (20), we get that, for each eigenvalue $\grave{\lambda}_k$, there exists at least one i such that

$$|\grave{\lambda}_k - \kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i)| \leq \sqrt{\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) - \delta^2} \sum_{\substack{j=1 \\ j\neq i}}^{m} \sqrt{\kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j)}.$$

■

Lemma A.2: The eigenvalues of a Gram matrix associated to a γ-coherent dictionary of m atoms are bounded as follows:

$$r^2 - (m-1)\gamma R^2 \leq \grave{\lambda}_m \leq \cdots \leq \grave{\lambda}_1 \leq R^2 + (m-1)\gamma R^2.$$

Proof: A γ-coherent dictionary satisfies

$$\max_{\substack{j=1\cdots m \\ j\neq i}} \frac{|\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)|}{\sqrt{\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i)\,\kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j)}} \leq \gamma,$$

for any $i, j = 1, 2, \ldots, m$, which yields

$$\begin{aligned}\max_{\substack{j=1\cdots m \\ j\neq i}} |\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)| &\leq \gamma \max_{\substack{j=1\cdots m \\ j\neq i}} \sqrt{\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i)\,\kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j)} \\ &= \gamma\sqrt{\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i)} \max_{\substack{j=1\cdots m \\ j\neq i}} \sqrt{\kappa(\grave{\boldsymbol{x}}_j, \grave{\boldsymbol{x}}_j)} \\ &\leq \gamma R\sqrt{\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i)}.\end{aligned}$$

By injecting this expression in (20), we get

$$\begin{aligned}\sum_{\substack{j=1 \\ j\neq i}}^{m} |\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)| &\leq (m-1) \max_{\substack{j=1\cdots m \\ j\neq i}} |\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)| \\ &\leq (m-1)\gamma R\sqrt{\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i)}.\end{aligned}$$

Since $\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i) \leq R^2$, this completes the proof.

■

Lemma A.3: The eigenvalues of a Gram matrix associated to a γ-Babel dictionary are bounded as follows:

$$r^2 - \gamma \leq \grave{\lambda}_m \leq \cdots \leq \grave{\lambda}_1 \leq R^2 + \gamma.$$

Proof: The proof is straightforward from the Geršgorin Discs Theorem, since expression (20) becomes

$$|\grave{\lambda}_k - \kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_i)| \leq \sum_{\substack{j=1 \\ j \neq i}}^{m} |\kappa(\grave{\boldsymbol{x}}_i, \grave{\boldsymbol{x}}_j)| \leq \gamma,$$

where the last inequality is due to the Babel measure. ■

REFERENCES

[1] "Special issue on signal processing for big data," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 8, pp. 507–507, June 2014.

[2] "Special issue on signal processing for big data," *Signal Processing, IEEE Transactions on*, vol. 62, pp. 1899–1899, April 2014.

[3] G. Giannakis, F. Bach, R. Cendrillon, M. Mahoney, and J. Neville, "Signal processing for big data [from the guest editors]," *Signal Processing Magazine, IEEE*, vol. 31, pp. 15–16, Sept 2014.

[4] J. Platt, "A resource-allocating network for function interpolation," *Neural Comput.*, vol. 3, pp. 213–225, June 1991.

[5] V. N. Vapnik, *Statistical Learning Theory*. New York, NY, USA: Wiley, September 1998.

[6] C. E. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.

[7] G. bin Huang, P. Saratch, S. Member, and N. Sundararajan, "A generalized growing and pruning rbf (ggap-rbf) neural network for function approximation," *IEEE Transactions on Neural Networks*, vol. 16, pp. 57–67, 2005.

[8] P. Honeine, "Online kernel principal component analysis: a reduced-order model," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, pp. 1814–1826, September 2012.

[9] B. Schölkopf, R. Herbrich, and A. J. Smola, "A generalized representer theorem," in *Proc. 14th Annual Conference on Computational Learning Theory and 5th European Conference on Computational Learning Theory*, COLT/EuroCOLT, (London, UK), pp. 416–426, Springer-Verlag, 2001.

[10] L. Csató and M. Opper, "Sparse online gaussian processes," *Neural Computation*, vol. 14, pp. 641–668, 2002.

[11] Y. Engel, S. Mannor, and R. Meir, "The kernel recursive least squares algorithm," *IEEE Trans. Signal Processing*, vol. 52, no. 8, pp. 2275–2285, 2004.

[12] P. P. Pokharel, W. Liu, and J. C. Principe, "Kernel least mean square algorithm with constrained growth," *Signal Processing*, vol. 89, no. 3, pp. 257 – 265, 2009.

[13] G. S. Babu and S. Suresh, "Meta-cognitive rbf network and its projection based learning algorithm for classification problems," *Appl. Soft Comput.*, vol. 13, pp. 654–666, Jan. 2013.

[14] Y.-K. Yang, T.-Y. Sun, C.-L. Huo, Y.-H. Yu, C.-C. Liu, and C.-H. Tsai, "A novel self-constructing radial basis function neural-fuzzy system," *Applied Soft Computing*, vol. 13, no. 5, pp. 2390 – 2404, 2013.

[15] P. Honeine, C. Richard, and J. C. M. Bermudez, "On-line nonlinear sparse approximation of functions," in *Proc. IEEE International Symposium on Information Theory*, (Nice, France), pp. 956–960, June 2007.

[16] C. Richard, J. C. M. Bermudez, and P. Honeine, "Online prediction of time series data with kernels," *IEEE Transactions on Signal Processing*, vol. 57, pp. 1058–1067, March 2009.

[17] J. A. Tropp, "Greed is good: algorithmic results for sparse approximation," *IEEE Trans. Information Theory*, vol. 50, pp. 2231–2242, 2004.

[18] M. Elad, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer, 2010.

[19] H. Fan, Q. Song, and S. B. Shrestha, "Online learning with kernel regularized least mean square algorithms," *Knowledge-Based Systems*, vol. 59, no. 0, pp. 21 – 32, 2014.

[20] W. Liu, J. C. Principe, and S. Haykin, *Kernel Adaptive Filtering: A Comprehensive Introduction*. Wiley Publishing, 1st ed., 2010.

[21] Z. Noumir, P. Honeine, and C. Richard, "One-class machines based on the coherence criterion," in *Proc. IEEE workshop on Statistical Signal Processing*, (Ann Arbor, Michigan, USA), pp. 600–603, 5–8 August 2012.

[22] Z. Noumir, P. Honeine, and C. Richard, "On simple one-class classification methods," in *Proc. IEEE International Symposium on Information Theory*, (MIT, Cambridge (MA), USA), pp. 2022–2026, 1–6 July 2012.

[23] R. Jenssen, "Mean vector component analysis for visualization and clustering of nonnegative data," *Neural Networks and Learning Systems, IEEE Transactions on*, vol. 24, pp. 1553–1564, Oct 2013.

[24] B. Schölkopf, A. Smola, and K.-R. Müller, "Nonlinear component analysis as a kernel eigenvalue problem," *Neural Comput.*, vol. 10, pp. 1299–1319, July 1998.

[25] C. Saidé, R. Lengellé, P. Honeine, and R. Achkar, "Online kernel adaptive algorithms with dictionary adaptation for mimo models," *IEEE Signal Processing Letters*, vol. 20, pp. 535–538, May 2013.

[26] G. Baudat and F. Anouar, "Kernel-based methods and function approximation," in *In International Joint Conference on Neural Networks (IJCNN)*, vol. 5, (Washington, DC, USA), pp. 1244–1249, July 2001.

[27] L. Csató and M. Opper, "Sparse representation for gaussian process models," in *Advances in Neural Information Processing Systems 13*, pp. 444–450, MIT Press, 2001.

[28] D. Nguyen-Tuong and J. Peters, "Incremental online sparsification for model learning in real-time robot control," *Neurocomputing*, vol. 74, no. 11, pp. 1859 – 1867, 2011.

[29] R. Rosipal, M. Koska, and I. Farkas, "Prediction of chaotic time-series with a resource-allocating RBF network," in *Neural Processing Letters*, pp. 185–197, 1997.

[30] B. Chen, S. Zhao, P. Zhu, and J. Principe, "Quantized kernel least mean square algorithm," *Neural Networks and Learning Systems, IEEE Transactions on*, vol. 23, pp. 22–32, Jan 2012.

[31] A. C. Gilbert, S. Muthukrishnan, M. J. Strauss, and J. Tropp, "Improved sparse approximation over quasi-incoherent dictionaries," in *International Conference on Image Processing (ICIP)*, vol. 1, (Barcelona, Spain), pp. 37–40, Sept. 2003.

[32] P. Honeine, "Analyzing sparse dictionaries for online learning with kernels," *IEEE Transactions on Signal Processing*, 2014 submitted.

[33] Z. Noumir, P. Honeine, and C. Richard, "Online one-class machines based on the coherence criterion," in *Proc. 20th European Conference on Signal Processing*, (Bucharest, Romania), pp. 664–668, 27–31 August 2012.

[34] P. Honeine, "An eigenanalysis of data centering in machine learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2014 submitted.

[35] R. A. Horn and C. R. Johnson, *Matrix analysis*. New York, NY, USA: Cambridge University Press, 2nd edition ed., December 2012.

PLACE PHOTO HERE

Paul Honeine (M'07) was born in Beirut, Lebanon, on October 2, 1977. He received the Dipl.-Ing. degree in mechanical engineering in 2002 and the M.Sc. degree in industrial control in 2003, both from the Faculty of Engineering, the Lebanese University, Lebanon. In 2007, he received the Ph.D. degree in Systems Optimisation and Security from the University of Technology of Troyes, France, and was a Postdoctoral Research associate with the Systems Modeling and Dependability Laboratory, from 2007 to 2008. Since September 2008, he has been an assistant Professor at the University of Technology of Troyes, France. His research interests include nonstationary signal analysis and classification, nonlinear and statistical signal processing, sparse representations, machine learning. Of particular interest are applications to (wireless) sensor networks, biomedical signal processing, hyperspectral imagery and nonlinear adaptive system identification. He is the co-author (with C. Richard) of the 2009 Best Paper Award at the IEEE Workshop on Machine Learning for Signal Processing. Over the past 5 years, he has published more than 100 peer-reviewed papers.